arXiv:1603.02740v2 [stat.ML] 13 Jun 2016

Pairwise Choice Markov Chains

Stephen Ragain, Johan Ugander
{sragain,jugander}@stanford.edu
Management Science & Engineering, Stanford University

Abstract

As datasets capturing human choices grow in richness and scale—particularly in online domains—there is an increasing need for choice models that escape traditional choice-theoretic axioms such as regularity, stochastic transitivity, and Luce's choice axiom. In this work we introduce the Pairwise Choice Markov Chain (PCMC) model of discrete choice, an inferentially tractable model that does not assume any of the above axioms while still satisfying the foundational axiom of *uniform expansion*, a considerably weaker assumption than Luce's choice axiom. We show that the PCMC model significantly outperforms the Multinomial Logit (MNL) model in prediction tasks on both synthetic and empirical datasets known to exhibit violations of Luce's axiom. Our analysis also synthesizes several recent observations connecting the Multinomial Logit model and Markov chains; the PCMC model retains the Multinomial Logit model as a special case.

1 Introduction

Discrete choice models describe and predict decisions between distinct alternatives. Classic applications include consumer purchasing decisions, choices of schooling or employment, and commuter choices for modes of transportation among available options. Early models of probabilistic discrete choice, including the well known Thurstone's Case V model [27] and Bradley-Terry-Luce (BTL) model [7], were developed and refined under diverse strict assumptions about human decision making. As increasingly complex individual choices become mediated by engineered and learned platforms—from online shopping to web browser clicking to interactions with recommendation systems—there is a pressing need for flexible models capable of describing and predicting nuanced choice behavior.

Luce's choice axiom, popularly known as the *independence of irrelevant alternatives* (IIA), is arguably the most storied assumption in choice theory [18]. The axiom consists of two statements, applied to any set of alternatives S within a broader universe U of alternatives. First, if $\Pr(a \text{ chosen from } \{a, b\}) = 0$ then it follows that $\Pr(a \text{ chosen from } S) = 0$ for all S containing a and b. Second, the probability that a is chosen from U conditioned on the chosen item lying in $S \subseteq U$ is precisely the probability that a is chosen from S. The BTL model, which defines $p_{ij} = \gamma_i/(\gamma_i + \gamma_j)$ for latent quality parameters $\gamma_i \geq 0$, satisfies the axiom while Thurstone's Case V model does not [1].

Soon after its introduction, the Bradley-Terry-Luce model was generalized from pairwise choices to choices from larger sets [4]. The resulting Multinomal Logit

(MNL) model again employs quality parameters $\gamma_i \geq 0$ for each $i \in U$ and defines p_{iS}, the probability of choosing i from $S \subseteq U$, proportional to γ_i for all $i \in S$. Any model satisfying Luce's choice axiom is equivalent to some MNL model [19].

One consequence of Luce's choice axiom, also shared by Thurstone's model, is *strict stochastic transitivity* between alternatives: if p_{xy} is the probability of choosing x over y among pairs $\{x, y\}$ in a universe U, then $p_{ab} \geq 0.5$ and $p_{bc} \geq 0.5$ imply $p_{ac} \geq \max(p_{ab}, p_{bc})$. A possibly undesirable consequence of strict stochastic transitivity is the necessity of a total order on the elements in U.

Many choice theorists and empiricists, including Luce, have noted that the choice axiom and stochastic transitivity are strong assumptions that do not hold for empirical choice data [9, 13, 14, 26, 28]. A range of discrete choice models striving to escape the confines of the choice axiom have emerged over the years. The most popular of these models have been Elimination by Aspects [29], the mixed MNL model [6], and the nested MNL model [22]. Inference is practically difficult for all three of these models [16, 23]. Additionally, Elimination by Aspects and the mixed MNL model also both exhibit the rigid property of regularity, defined below.

A broad, important model class in the study of discrete choice is the class of *random utility models* (RUMs) [4, 20]. A RUM affiliates with each $i \in U$ a random variable X_i and defines for each subset S of U the probability $\Pr(i \text{ chosen from } S) = \Pr(X_i \geq X_j, \forall j \in S)$. An *independent RUM* has pairwise independent X_i. RUMs in general assume regularity but do not necessarily assume the choice axiom or stochastic transitivity. Thurstone's Case V model and the BTL model are both independent RUMs; the mixed MNL model and the Elimination by Aspects model are both RUMs. The nested MNL model, however, is not a RUM. A major result by McFadden and Train establishes that for any RUM there exists a mixed MNL model that can approximate the choice probabilities of that RUM to within an arbitrary error [23], a strong result about the generality of mixed MNL models.

When considering models circumventing the choice axiom, the weaker property of *regularity* often holds: for any choice sets A, B where $A \subseteq B$, $\Pr(x \text{ chosen from } A) \geq \Pr(x \text{ chosen from } B)$. For example, all RUMs exhibit regularity, even if they don't exhibit the choice axiom. Regularity, though intuitively pleasing, has empirical drawbacks: it prevents models from expressing framing effects [13] and other observations from modern behavior economics [28]. This rigidity motivates us to contribute a new model of discrete choice that escapes historically common assumptions while still furnishing enough structure to be inferentially tractable.

The present work. In this work we introduce a simple, inferentially tractable model of discrete choice that we call the Pairwise Choice Markov Chain (PCMC) model. The PCMC model is based on a continuous time Markov chain (CTMC) on the subsets S of the universe U. The transition rate parameters between alternatives are given by a matrix P of pairwise selection probabilities. The model conceives of a selection probability p_{iS} as the probability mass on alternative $i \in S$ of the stationary distribution of the CTMC on S.

The transition rates of the CTMC can be interpreted as a simple transformation of the parameter matrix P of pairwise selection probabilities, and as such the PCMC model can be thought of as defining set-wise selection probabilities strictly in terms of pairwise selection probabilities. One of the strengths of the PCMC model is that this

matrix can be parameterized further by arbitrary models of pairwise choice. Indeed, we show that when the matrix P is parameterized as the pairwise selection probabilities of a BTL pairwise model, the PCMC model reduces to an MNL model. Recent parametrization of non-transitive pairwise probabilities such as the Blade-Chest model [8] can be usefully employed to reduce the number of free parameters of the PCMC model.

Our PCMC model can be thought of as building upon the observation underlying the recently introduced Iterative Luce Spectral Ranking (I-LSR) procedure for efficiently finding the maximum likelihood estimate for parameters of MNL models [21]. The analysis of I-LSR is precisely analyzing a PCMC model in the special case where the matrix P has been parameterized by BTL. In that case the stationary distribution of the chain is found to match the stationary conditions of the MNL likelihood function, establishing a deep connection between MNL models and Markov chains. The PCMC model generalizes that connection.

The analysis of I-LSR is one of several recent works to realize connections between the MNL model and Markov chains. Other recent connections have been made by the work on RankCentrality [24], which employs a discrete time Markov chain for inference in the place of I-LSR's continuous time chain. The RankCentrality chain can be interpreted as a δ-skeleton discretization of our CTMC. Separate work has contributed a different discrete time Markov chain model of choice substitution capable of approximating any RUM [3], a related problem but one with a strong focus on ordered preferences.

The PCMC model is not a RUM, and in general does not exhibit stochastic transitivity, regularity, or obey the choice axiom. We find that the PCMC model does, however, obey the lesser known but fundamental axiom of *uniform expansion*, a weakened version of Luce's choice axiom proposed by Yellott that implies the choice axiom for models that are independent RUMs [30]. In this work we define a convenient structural property termed contractibility, for which uniform expansion is a special case, and we show that the PCMC model exhibits contractibility. Of the models mentioned above, only Elimination by Aspects exhibits uniform expansion without being an independent RUM. Elimination by Aspects obeys regularity, which the PCMC model does not; as such, the PCMC model is uniquely positioned in the literature of axiomatic discrete choice, minimally satisfying uniform expansion without the other aforementioned axioms.

After presenting the model and its properties, we investigate choice predictions from our model on diverse synthetic datasets as well as two empirical choice datasets. The empirical choice datasets concern transportation choices made on commuting and shopping trips in San Francisco. Meanwhile for synthetic data, when data is generated from a PCMC model exhibiting violations of Luce's choice axiom then a PCMC model outperforms an MNL model on inference, and a PCMC model is competitive with an MNL model on data consistent with Luce's choice axiom. More significantly, for both of the empirical datasets we find that a learned PCMC model predicts empirical choices significantly better than a learned MNL model.

2 The PCMC model

The Pairwise Choice Markov Chain (PCMC) models the selection probability p_{iS}, the probability of choosing i from $S \subseteq U$, as the probability mass on alternative $i \in S$ of the stationary distribution of a continuous time Markov chain (CTMC) on the set of alternatives S. The transition rates of the chain are given as a simple linear function of a parameter matrix P that is itself interpretable as a matrix of pairwise selection probabilities.

Formally, given $S \subseteq U$ we construct the rate matrix $Q_S = \{q_{ij}\}_{i,j \in S}$ as follows:

$$q_{ij} = p_{ji}, \text{ for } i \neq j, \qquad q_{ii} = -\sum_{j \neq i} q_{ij}.$$

Let π_S be the stationary distribution of the corresponding CTMC on S. We define $p_{iS} := \pi_S(i)$, and also occasionally abuse notation to write $\pi_S(A) = \sum_{x \in A} \pi_S(x)$ when the intention is clear. In a final additional abuse of notation, when $|S| = 2$ we write p_{ij} instead of p_{iS} or $p_{i\{i,j\}}$. This abuse is motivated in the next section, where we show that the PCMC model agrees with the input pairwise probabilities when queried about pairs, i.e. $\pi_{\{i,j\}}(i) = p_{ij}$.

We now show that the PCMC model is well defined.

Proposition 1. *The choice probabilities p_{iS} are well defined for all $i \in S$, all $S \subseteq U$ of a finite U.*

Proof. We will show that there exists a single closed communicating class, which implies a unique stationary distribution. A closed communicating class C is a non-empty subset of S such that the chain will never leave C once entering it, and each state in C is reachable from every other state in C.

Because S is finite, there must be at least one closed communicating class. Suppose the chain had more than one closed communicating class and that $i \in S$ and $j \in S$ were in different closed communicating classes. But $p_{ij} + p_{ji} = 1$, so at least one of q_{ji} and q_{ij} is strictly positive. But then the chain can switch communicating classes through the strictly positive edge, a contradiction. □

Note that the support of π_S is the closed communicating class, but S may have transient states corresponding to alternatives with a selection probability of zero.

In the subsections that follow, we develop key properties of the model. We first show that pairwise choice probabilities under the PCMC model are precisely the entries of the parameter matrix P. We then show and give an example of a PCMC model which is neither regular nor a RUM. Lastly we show that a PCMC model with pairwise probabilities originating from a Bradley-Terry-Luce (BTL) pairwise model is equivalent to a Multinomial Logit (MNL) set-wise model, and that the PCMC model can be used as a wrapper around other models for pairwise choice.

2.1 Preservation of pairwise probabilities

Here we demonstrate that the PCMC model amounts to an identity mapping on pairs $S = \{i, j\}$, alleviating any confusion about the term "pairwise probabilities," as the

pairwise predictions of the model agrees with its own parameters. Consider a pair set $S = \{i, j\}$. When $p_{ij} \in (0, 1)$, then the unique stationary distribution π_S satisfying $\pi_S Q_S = 0$ subject to $||\pi_S||_1 = 1$ will have

$$\pi_S(j)q_{ji} = \pi_S(i)q_{ij},$$
$$\pi_S(i) + \pi_S(j) = 1.$$

These equations yield $\pi_S(i) = p_{ij}, \pi_S(j) = p_{ji}$. The corner cases where $p_{ij} = 0$ or $p_{ij} = 1$ have the dominating state as an absorbing state, also yielding agreement with π_S. We conclude that when the PCMC model is asked to output selection probabilities on sets of size two, the output is consistent with the input pairwise probability parameters.

2.2 Regularity and RUMs

The regularity property requires that the probability of choosing $a \in S$ does not decrease when removing elements from S. More compactly, $p_{aS'} \geq p_{aS}$ for all $S' \subseteq S \subseteq U$. We will now show that the PCMC model is not constrained to exhibit regularity, nor is it a random utility model (RUM) in general. Note that RUMs exhibit regularity because $P(X_i = \max_{j \in S'} X_j) \geq P(X_i = \max_{j \in S} X_j)$ for all $S' \subseteq S \subseteq U$.

Consider $U = \{r, p, s\}$ corresponding to a stochastic rock-paper-scissors-like game, in the sense that

$$P = \begin{bmatrix} 0 & 1-\alpha & \alpha \\ \alpha & 0 & 1-\alpha \\ 1-\alpha & \alpha & 0 \end{bmatrix}$$

for some $\alpha \in [0, 1]$, where the rows of P are indexed by r, p, and s respectively. For any pair set we have $p_{ij} = \alpha$ if i beats j in the game. Meanwhile, the PCMC model stipulates $p_{rU} = p_{pU} = p_{sU} = 1/3$ regardless of α. Thus, regularity does not hold for $\alpha < 1/3$ or $\alpha > 2/3$.

The PCMC model's lack of regularity is a positive trait in the sense that empirical choice phenomena such as framing effects and asymmetric dominance violate regularity [13], and the PCMC model is rare in its ability to model such choices. Deriving necessary and sufficient conditions on P for a PCMC model to be a RUM analogous to known characterization theorems for RUMs [10] and known sufficient conditions for nested MNL models to be RUMs [5] is an interesting open challenge.

2.3 Multinomial Logit from Bradley-Terry-Luce

We now observe that the Multinomial Logit Model (MNL), also called the Plackett-Luce model, is precisely a PCMC model with a matrix P consisting of pairwise BTL probabilities. Recall that the BTL model assumes the existence of "quality scores" $\gamma_i > 0$ for $i \in U$ with $p_{ij} = \gamma_i/(\gamma_i + \gamma_j), \forall i, j \in U$ and that the MNL generalization defines $p_{iS} \propto \gamma_i, \forall i \in S$ for each $S \subseteq U$.

Proposition 2. *If the pairwise probabilities P are derived from a BTL model with parameters γ, then for each $S \subseteq U$, the PCMC probabilities p_{iS} are consistent with an MNL model on S with parameters γ.*

Proof. Let the PCMC rate matrix Q_S be constructed from a matrix of pairwise probabilities P generated from a BTL model with parameters γ. We aim to show that $\pi_S = \frac{\gamma}{||\gamma||_1}$ is a stationary distribution of the PCMC chain: $\pi_S^T Q_S = 0$. We have:

$$\begin{aligned}(\pi_S^T Q_S)_i &= \frac{1}{||\gamma||_1}\left(\sum_{j\neq i}\gamma_j p_{ij} - \gamma_i(\sum_{j\neq i} p_{ij})\right) \\ &= \frac{\gamma_i}{||\gamma||_1}\left(\sum_{j\neq i}\frac{\gamma_j}{\gamma_i+\gamma_j} - \sum_{j\neq i}\frac{\gamma_j}{\gamma_i+\gamma_j}\right) = 0,\ \forall i.\end{aligned}$$

Thus π_S is always the stationary distribution of the chain, and we know by Proposition 1 that it is unique. It follows that $p_{iS} \propto \gamma_i$ for all $i \in S$, as desired. □

The observation in this subsection also helps explain why the recently introduced Iterative Luce Spectral Ranking (I-LSR) procedure succeeds in inferring the parameters of an MNL model from the stationary distribution of an appropriately constructed continuous-time Markov chain. If the choice data being subjected to inference is pairwise, the chain that the I-LSR procedure constructs is precisely a PCMC chain with P parameterized by iteratively inferred BTL pairwise choice probabilities.

2.4 Parameterizations of the P matrix

Any model for pairwise probabilities can be used to extend the PCMC model by parametrizing $P = P(\theta)$ for some parameters θ. We saw above that when using a BTL model to parameterize P with quality scores γ, the resulting PCMC model on $P(\gamma)$ was equivalent to MNL. We demonstrate this capacity for general parameterization with the recently introduced Blade-Chest model of non-transitive pairwise choices.

Blade-Chest models are based on geometric d-dimensional embeddings of alternatives, where each alternative i is parameterized with a blade vector $b_i \in \mathbb{R}^d$ and a chest vector $c_i \in \mathbb{R}^d$, in addition to a BTL-like quality parameter $\gamma_i > 0$. The Blade-Chest model comes in two variations: the *Blade-Chest distance model*, where

$$p_{ij}(b, c, \gamma) = S(||b_i - c_j||_2^2 - ||b_j - c_i||_2^2 + \gamma_i - \gamma_j),$$

where $S(x) = (1+\exp(-x))^{-1}$ is the sigmoid/logistic function, and the *Blade-Chest inner product model*, where

$$p_{ij}(b, c, \gamma) = S(b_i \cdot c_j - b_j \cdot c_i + \gamma_i - \gamma_j).$$

The quality parameters γ_i serve to connect the models to the BTL model, but do not meaningfully increase their expressiveness, so we disregard them in our use of the Blade-Chest distance model in this work. Thus, these two Blade-Chest models provide

useful parameterizations of non-transitive pairwise probability matrices, $P(\theta)$, with $\theta = \{b, c\}$ consisting of the $2dn$ parameters of the blade and chest embeddings.

In Section 4 we demonstrate on synthetic data that a PCMC model wrapped around a Blade-Chest model maintains the performance of a PCMC model while significantly reducing the number of parameters.

Another technique for parametrizing P involves representing p_{ij} as a function of features of i and j, i.e. $p_{ij} = f(X_i, X_j; \theta)$ where X_i gives the salient features of i, and θ represents parameters that, for instance, give weights to these features. We can also formulate such analysis as a factoring $P = W^T F$ where W is a weight matrix and F a feature matrix. We do not explore such parameterizations in this work, but merely highlight the potential to employ latent features of objects in a straight-forward manner, an approach closely related to conjoint analysis [11].

3 Properties

While we have demonstrated already that the PCMC model avoids several restrictive properties that are often inconsistent with empirical choice data, we demonstrate in this section that the PCMC model exhibits deep structure in the form of contractibility, which implies uniform expansion. Inspired by a thought experiment that was posed as an early challenge to the choice axiom, we define the property of contractibility to handle notions of similarity between elements. We demonstrate that the PCMC model has contractibility, and gracefully handles this thought experiment.

3.1 Uniform expansion

Yellott [30] introduced *uniform expansion* as a weaker condition than Luce's choice axiom, but one that implies the choice axiom in the context of any independent random utility model. Yellott posed the axiom of invariance to uniform expansion in the context of "copies" of elements.

Definition 1 (Copies)**.** *For i, j in $S \subseteq U$, we say that i and j are copies if for all $k \in S - i - j$, $p_{ik} = p_{jk}$ and $p_{ij} = 1/2$.*

Yellott's introduction to uniform expansion asks the reader to consider an offer of a choice of beverage from k identical cups of coffee, k identical cups of tea, and k identical glasses of milk. Yellott contends that probability the person chooses a type of beverage in this scenario should be the same as if they were only shown one beverage of each type.

Definition 2 (Uniform Expansion)**.** *Consider a choice between n elements in a set $S_1 = \{i_{11}, \ldots, i_{n1}\}$, and another choice from a set S_k containing k copies of each of the n elements: $S_k = \{i_{11}, \ldots, i_{1k}, i_{21}, \ldots, i_{2k}, \ldots, i_{n1}, \ldots, i_{nk}\}$. The axiom of uniform expansion states that for each $m = 1, \ldots, n$:*

$$p_{i_{m1}S_1} = \sum_{j=1}^{k} p_{i_{mj}S_k}.$$

We will show that the PCMC model always exhibits a more general property of contractibility, of which uniform expansion is a special case; it thus always exhibits uniform expansion.

Yellott demonstrated that for any independent random utility model, the double-exponential distribution family is the only family of independent distributions that exhibit uniform expansion, and that Thurstone's model based on the Normal family in particular does not exhibit uniform expansion.

While the axiom feels natural as a model for discrete choice, it should be regarded with some skepticism in applications that model competitions. For example, in sports matches or races, the winner of the competition can be modeled as the competitor with highest skill on that day; competitions are often modeled using random utility models. If we have a competitor with a skill distribution whose wins come from occasional "good days" (their skill distribution has a heavy upper tail) then uniform expansion should not hold, and is not a desirable axiom. This example relates to recent work on team performance and selection [15], where non-invariance under uniform expansion plays a key role.

3.2 Contractibility

In a book review of Luce's early work on the choice axiom, Debreau [9] considers a choice between three musical recordings: one of Beethoven's eighth symphony conducted by X, another of Beethoven's eighth symphony conducted by Y, and one of Debussy quartet conducted by Z. We will call these options B_1, B_2, and D respectively. When compared to D, Debreau argues that B_1 and B_2 are indistinguishable in the sense that $p_{DB_1} = p_{DB_2}$. However, an agent can reasonably prefer B_1 over B_2 in the sense that $p_{B_1B_2} > 0.5$, bringing us outside the realm of copies.

Debreu's thought experiment is intended to show a weakness of the BTL model: $p_{DB_1} = p_{DB_2}$ implies that $\gamma_{B_1} = \gamma_{B_2}$, in turn implying that $p_{B_1B_2} = 0.5$, so the above choice can not be encoded by any model that employs Luce's axiom.

To address contexts in which elements compare identically to alternatives but not each other (e.g. B_1 and B_2), we introduce *contractible partitions* that group these similar alternatives into sets. We then show that when a PCMC model contains a contractible partition, the relative probabilities of selecting from one of these partitions is independent from the pairwise probability within the partition. Our contractible partition definition can be viewed as akin to (but distinct from) *nests* in nested MNL models [22].

Definition 3 (Contractible Partition)**.** *A partition of U into non-empty sets $A_1, \ldots, A_k$ is a contractible partition if $p_{a_i a_j} = \lambda_{ij}$ for all $a_i \in A_i, a_j \in A_j$ for some $\Lambda = \{\lambda_{ij}\}$ for $i, j \in \{1, \ldots, k\}$.*

Proposition 3. *For a given Λ, let $A_1, \ldots, A_k$ be a contractible partition for two PCMC models on U represented by P, P' with stationary distributions π, π'. Then for any A_i:*

$$\sum_{j \in A_i} p_{jU} = \sum_{j \in A_i} p'_{jU}, \tag{1}$$

or equivalently, $\pi(A_i) = \pi'(A_i)$.

Proof. Suppose P has contractible partition $A_1, \ldots, A_k$ with respect to Λ. Then $\pi^T Q_U = 0$ gives for $x \in A_1$ WLOG,

$$\pi(x)\left(\sum_{y\in A_1\setminus x} q_{xy} + \sum_{i=2}^{k}\sum_{a_i\in A_i} q_{xa_i}\right) = \sum_{y\in A_1\setminus x}\pi(y)q_{yx} + \sum_{i=2}^{k}\sum_{a_i\in A_i}\pi(a_i)q_{a_ix}. \quad (2)$$

Noting that for $a_i \in A_i$ and $a_j \in A_j$, $q_{a_ia_j} = p_{a_ja_i} = \lambda_{ji}$, and (2) can be rewritten:

$$\pi(x)\left(\sum_{y\in A_1\setminus x} q_{xy}\right) + \pi(x)\sum_{i=2}^{k}|A_i|\lambda_{i1} = \sum_{y\in A_1\setminus x}\pi(y)q_{yx} + \sum_{i=2}^{k}\pi(A_i)\lambda_{1i}.$$

Summing over $x \in A_1$ then gives

$$\sum_{x\in A_1}\pi(x)\left(\sum_{y\in A_1\setminus x} q_{xy}\right) + \pi(A_1)\sum_{i=2}^{k}|A_i|\lambda_{i1} =$$
$$\sum_{x\in A_1}\sum_{y\in A_1\setminus x}\pi(y)q_{yx} + |A_1|\sum_{i=2}^{k}\pi(A_i)\lambda_{1i}.$$

The leftmost term of each side is equal, so we have

$$\pi(A_1) = \frac{|A_1|\sum_{i=2}^{k}\pi(A_i)\lambda_{1i}}{\sum_{i=2}|A_i|\lambda_{i1}}, \quad (3)$$

which makes $\pi(A_1)$ the solution to global balance equations for a continuous time Markov chain with the states $\{A_1, \ldots, A_k\}$ and transition rate $\tilde{q}_{A_iA_j} = |A_j|\lambda_{ji}$ between state A_i and A_j, and $\tilde{q}_{A_iA_i} = -\sum_{j\neq i}\tilde{q}_{A_iA_j}$. Since the partition contains no empty sets, we can use $\lambda_{ij}+\lambda_{ji} = 1$ to show that the induced chain has a unique closed communicating class (as in the proof of Proposition 1). We conclude that $\{\pi(A_i)\}_{i=1}^{k}$ are all fixed if $A_1, ..., A_k$ is a contractible partition, as the transition rates of the new chain are determined entirely by Λ and $|A_1|, \ldots, |A_k|$. □

The intuition is that we can "contract" each A_i to a single "type" because the probability of choosing some $x \in A_i$ is independent of the pairwise probabilities between elements within the sets. The above proposition and the contractibility of a PCMC model on all uniformly expanded sets implies that all PCMC models exhibit uniform expansion.

Proposition 4. *The PCMC model exhibits uniform expansion.*

Proof. We translate the problem of uniform expansion into the language of contractibility. Let U_1 be the universe of unique items $i_{11}, i_{21}, \ldots, i_{n1}$ and let U_k be a universe containing k copies of each item in U_1. Let i_{mj} denote the jth copy of the mth item in U_1. Thus $U_k = \cup_{m=1}^{n}\cup_{j=1}^{k} i_{mj}$.

Let Q be the transition rate matrix of the CTMC on U_1. We construct a contractible partition of U_k into the n sets, each containing the k copies of some item in U_1. Thus

$A_m = \cup_{j=1}^{k} i_{mj}$. This means that $\{A_m\}_{m=1}^{n}$ is a contractible partition of U_k with $\Lambda = Q$. Noting $|A_m| = k$ for all m in Equation (3) above results in $\{\pi(A_m)\}_{m=1}^{n}$ being the solution to $\pi^T Q = \pi^T \Lambda = 0$. Thus $p_{i_m U_1} = \pi(A_m) = \sum_{j=1}^{k} p_{i_{mj} U_k}$ for each m. □

It is clear that uniform expansion would hold even if the pairwise probabilities between copies were not 1/2. Note also that every PCMC model on n elements contains a trivial contractible partition $A_1 = \{x_1\}, ..., A_n = \{x_n\}$. It is an interesting open question wether one could perhaps parameterize P directly in terms of contractible partitions.

4 Inference and prediction

Our goal in formulating this model is ultimately to use past choices from diverse subsets $S \subseteq U$ to predict future choices. In this section we first give the log-likelihood function $\log \mathcal{L}(P; \mathcal{C})$ of the PCMC matrix P given a choice data collection of the form $\mathcal{C} = \{(i_k, S_k)\}_{k=1}^{n}$, where $i_k \in S_k$ was the item chosen from S_k. Let $C_{iS}(\mathcal{C}) = \#\{(i_k, S_k) \in \mathcal{C} : i_k = i, S_k = S\}$ denote the number of times in the data that i was chosen out of set S for each $S \subseteq U$, and let $C_S(\mathcal{C}) = \#\{(i_k, S_k) \in \mathcal{C} : S_k = S\}$ be the number of times that S was the choice set for each $S \subseteq U$. We then investigate the ability of a learned PCMC model to make choice predictions on empirical data, benchmarked against a learned MNL model, and interpret the inferred parameters $\hat{P}$. We conclude by examining the performance of two leaned PCMC models, one wrapped around a Blade-Chest model, and an MNL model on three synthetic datasets.

4.1 Maximum likelihood

For each $S \subseteq U$, $i \in S$, recall that $p_{iS}(P)$ is the probability that i is selected from set S as a function of the parameter matrix P. After dropping all additive constants, the log-likelihood of P given the data $\mathcal{C}$ (derived from the probability mass function of the multinomial distribution) is:

$$\log \mathcal{L}(P; \mathcal{C}) = \sum_{S \subseteq U} \sum_{i \in S} C_{iS}(\mathcal{C}) \log(p_{iS}(P)).$$

Recall that for the PCMC model, $p_{iS}(P) = \pi_S$, where π_S is the stationary distribution for a PCMC chain with rate matrix $Q_S(P)$, i.e. π_S can be computed by solving $\pi_S^T Q_S(P) = 0$ subject to the constraint $\sum_{i \in S} \pi_S(i) = 1$. Thus, $p_{iS}(P)$ is an implicit function of those elements p_{ij} of P where $i, j \in S$, but the definition in terms of the stationary distribution means that there is no closed form expression for $p_{iS}(P)$.The implicit definition also makes it difficult to derive gradients for $\log \mathcal{L}$ with respect to some p_{ij}. We employ BFGS [25] to maximize $\log \mathcal{L}(P; \mathcal{C})$, which is non-concave in general.

When the pairwise probability matrix P is parameterized by the Bradley-Terry model, $P(\gamma)$ with $p_{ij} = \frac{\gamma_i}{\gamma_i + \gamma_j}$, we have shown earlier that our PCMC model is equivalent to MNL, and equivalently the likelihood function reduces to that of MNL. In

these cases one can use Iterative Luce Spectral Ranking (I-LSR), a recent efficient algorithm for inferring MNL parameters γ, which has been shown to be much faster than traditional MM or Newton methods for inferring MNL parameters [21].

4.2 Empirical data results

We evaluate our inference procedure on two empirical choice datasets, `SFwork` and `SFshop`, collected from a survey of transportation preferences around the San Francisco Bay Area [17]. The `SFshop` dataset contains 3,157 observations each consisting of a choice set of transportation alternatives available to individuals traveling to and returning from a shopping center, as well as a choice from that choice set. The `SFwork` dataset, meanwhile, contains 5,029 observations consisting of commuting options and the choice made on a given commute. Basic statistics describing the choice set sizes and the number of times all pairs of alternatives appear in the same choice set appear in the appendix.

The majority of datasets in the literature on discrete choice focus on pairwise comparisons or ranked lists, where lists inherently assume transitivity and the independence of irrelevant alternatives. The `SFwork` and `SFshop` datasets are rare examples of public datasets that genuinely study choices from sets larger than pairs. Additional evaluations of predictive performance across a range of synthetic datasets follows this empirical evaluation.

We train our model on observations $T_{\text{train}} \subset \mathcal{C}$ and evaluate on a test set $T_{\text{test}} \subset \mathcal{C}$ via

$$\text{Error}(T_{\text{train}}; T_{\text{test}}) = \frac{1}{|T_{\text{test}}|} \sum_{(i,S) \in T_{\text{test}}} \sum_{j \in S} |p_{jS}(\hat{P}(T_{\text{train}})) - \tilde{p}_{iS}(T_{\text{test}})|, \tag{4}$$

where $\hat{P}(T_{\text{train}})$ is the estimate for P obtained from the observations in T_{train} and $\tilde{p}_{iS}(T_{\text{test}}) = C_{iS}(T_{\text{test}})/C_S(T_{\text{test}})$ is the empirical probability of i was selected from S among observations in T_{test}. Note that $\text{Error}(T_{\text{train}}; T_{\text{test}})$ is the expected ℓ_1-norm of the difference between the empirical distribution and the inferred distribution on a choice set drawn uniform at random from the observations in T_{test}.

Figure 1 shows $\text{Error}(T_{\text{train}}; T_{\text{test}})$ on the `SFshop` and `SFwork` data as the learning procedure is applied to increasing amounts of data. The results are averaged over 1,000 permutations σ of the data such that for a given permutation $T_{\text{train}} = \{(i_k, S_k) : \sigma(k) \leq \frac{m}{10} n\}$ and $T_{\text{test}} = \{(i_k, S_k) : \frac{m}{10} n < \sigma(k) \leq \frac{m+1}{10} n\}$ for $m = 1, \ldots, 9$. More simply, we train on the first m tenth of shuffled data and test on the $(m+1)$th tenth, averaging our results from 1000 uniformly random shuffles. We see that our model is better equipped to learn from and make predictions in both datasets, and observe a 7.3% and 16.1% reduction in the average error on the two datasets.

Figure 2 gives heat maps of $\hat{P}_{\text{work}}$ and $\hat{P}_{\text{shop}}$, the pairwise probabilities inferred by the PCMC model from the `SFwork` and `SFshop` datasets respectively. The index ordering of each matrix follows the estimated selection probabilities (the stationary distribution) of the PCMC model on the full set of the alternatives for that dataset. Numerical values for the entries of $\hat{P}_{\text{work}}$ and $\hat{P}_{\text{shop}}$ appear in the appendix.

The ordered options for `SFwork` are: (1) driving alone, (2) walking, (3) sharing a ride with one other person, (4) public transit, (5) biking, and (6) carpooling with at

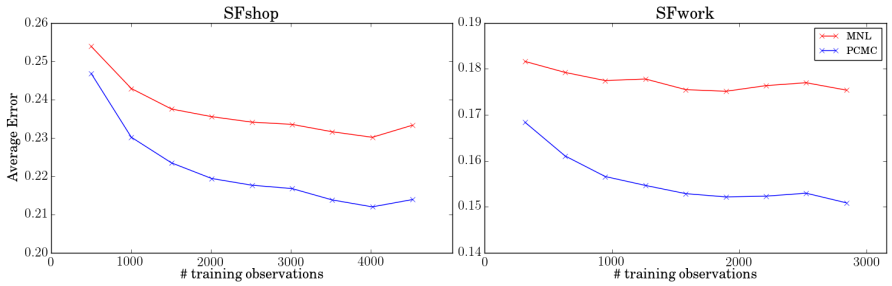

Figure 1: Prediction error on `SFshop` and `SFwork` data for the PCMC and MNL models, The mean improvement in the prediction error when training on 90% of the data is 7.3% for `SFshop` and 16.1% for `SFwork`.

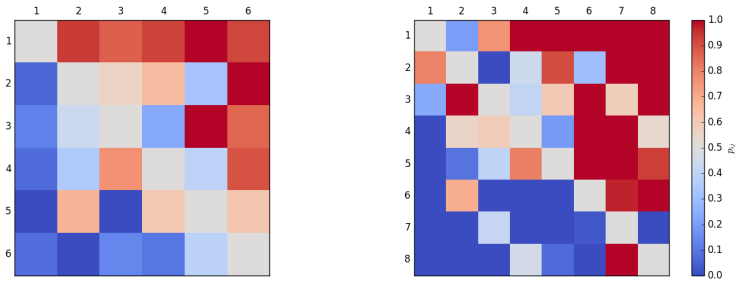

Figure 2: Heat maps of the $\hat{P}_{\text{work}}$(left) and $\hat{P}_{\text{shop}}$ (right) matricies inferred when training a PCMC model on `SFwork` and `SFshop` respectively.

least two others. The ordered options for `SFshop` are (1) drive alone both directions, (2) share a ride with one person in one direction and drive alone in the other direction, (3) share a ride with one person in both directions, (4) share a ride with more than one person in both directions, (5) walk, (6) share a ride with one person in one direction and more than one in the other direction, (7) bike, and (8) take public transit.

In the $\hat{P}$ inferred from `SFshop` data, we see that walking (option 2) has the second highest selection probability from the universe and biking (option 5) the second lowest, but the inferred pairwise probability in the PCMC model predicts biking would often be chosen over walking when they are the only two alternatives. Of the twenty choice triplets in the `SFwork` data, two triplets exhibit $\hat{p}_{ij}$, $\hat{p}_{jk}$, and $\hat{p}_{ki}$ all greater than 0.5: (4, 3, 2) where the minimum probability is 0.607 and (4, 2, 1) where the minimum probability is 0.559. This demonstrates that the PCMC model has inferred stochastic intransitivity among the alternatives.

Constructing a tournament graph on the alternatives where $(i, j) \in E$ if $p_{ij} \geq 0.5$, transitive triplets are then length-3 cycles in the tournament. A bound due to Harary and Moser [12] establishes that the maximum number of cyclic triples on a tournament graph on n nodes as 8 when $n = 6$ and 20 when $n = 8$. We conclude that according to our learned model the choices exhibit two out of a maximum 8 possible cyclic triplets in the `SFwork` data, and 6 out of 20 possible cyclic triplets for the `SFshop` data.

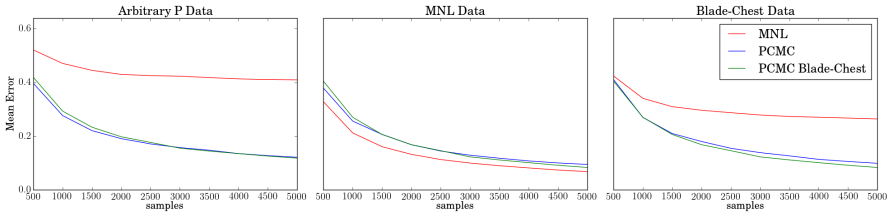

Figure 3: Learning error for an arbitrary PCMC model, MNL model using I-LSR, and Blade-Chest PCMC model on synthetic data generated from (i) Arbitrary P, (ii) MNL, and (iii) Blade-Chest parameterizations.

4.3 Synthetic data results

We further evaluate our inference procedure's performance in three synthetic data regimes: (i) choice data generated from a PCMC model with an "arbitrary" P ($p_{ij} \sim \text{UNIF}[0,1]$ i.i.d. for $i < j$, $p_{ji} = 1 - p_{ij}$), (ii) choice data generated for a simple MNL model with qualities $\gamma_i \sim \text{UNIF}[0,1]$, and (iii) choice data generated from a PCMC model with a pairwise matrix P parameterized by a two-dimensional Blade-Chest distance model, as discussed previously in Section 2.4. In order to create a strongly non-transitive instance of the Blade-Chest distance model, we draw the blades b_i and chests c_i uniformly at i.i.d. points along the two-dimensional unit circle, naturally producing many triadic impasses. The PCMC model's P matrix has $\binom{n}{2}$ parameters in general. When P is parameterized according to BTL we have just n parameters, and when it is parameterized according to a Blade-Chest distance model in d dimensions, we have $2dn$ parameters.

We evaluate each parameterization (arbitrary, MNL, Blade-Chest distance) for each synthetic regime. We employ the Iterative Luce Spectral Ranking (I-LSR) algorithm to learn the model under the BTL parameterization of P, where the PCMC model is equivalent to an MNL model. When the data is generated by MNL, we expect an MNL model to outperform inference under the general parameterization. When the data is generated by a non-MNL PCMC model, we expect MNL to exhibit restricted performance compared to a general parameterization, since the data is not generated by a model consistent with Luce's choice axiom.

Training data T_{train} and test data T_{test} are generated according to each of the three PCMC models, and then learned for each of the three parameterizations. For each model we draw a stream of triplets S uniformly at random from the size three subsets of U. We then infer each model, and examine the evaluation data using the same error metric as for the empirical data.

Figure 3 shows our error as the data grows, averaged across 10 instances, for each data generating process and each inference parameterization. First, each model is trained on an episode of 500 samples, and error is measured on predictions for the next episode of 500 samples. These new samples are then added to the training data, and error is again measured on a new episode of 500 samples, for each inference scheme.

The inference is applied to a set U with $n = 20$ objects, meaning that MNL has 20 parameters, the PCMC model with arbitrary P has $\binom{20}{2} = 190$ parameters, and

the PCMC model with the Blade-Chest distance parameterization in $\mathbb{R}^2$ uses $2dn = 80$ parameters. Overall we examine 9 data–model pairs, trained sequentially in 10 episodes, averaged across 10 instances. The figure thus represents 900 trained models.

As expected, the best PCMC model in each case is the model used to generate each dataset. More significantly, the Blade-Chest parametrization of the PCMC model performs very similarly to the general PCMC model in all three scenarios, despite having far fewer parameters. This is promising in domains where the $O(n^2)$ parameters of a general PCMC model is infeasible but the $O(n)$ parameterization used by Blade-Chest may be feasible.

5 Conclusion

We introduce a Pairwise Choice Markov Chain (PCMC) model of discrete choice capable of extending pairwise probabilities to larger sets. The PCMC model uses the stationary distribution of a continuous time Markov chain on the set of candidate choices to define selection probabilities, where the transition rates come from the pairwise selection probabilities.

In general the model is not a random utility model (RUM), and maintains broad flexibility by eschewing the implications of Luce's choice axiom, stochastic transitivity, and regularity. Despite this flexibility, we demonstrate that the PCMC model exhibits desirable structure by fulfilling uniform expansion, a property previously found only in the Multinomial Logit (MNL) model and the intractable Elimination by Aspects model.

We also introduced the notion of contractibility, a property underlying important thought experiments instrumental in moving choice theory beyond the choice axiom, and showed that the PCMC model exhibits contractibility. We showed that the PCMC model offers straightforward inference through maximum likelihood estimation, and that a learned PCMC model predicts empirical choice data with a significantly higher fidelity than a learned MNL model. We also demonstrated through synthetic data both that a PCMC model remains competitive with an MNL model on data consistent with Luce's choice axiom and that wrapping a PCMC model around a Blade-Chest model can significantly reduce the number of parameters without compromising inferential performance.

The flexibility and tractability of the PCMC model opens up many compelling research directions. First, what necessary and sufficient conditions on the matrix P guarantee that a PCMC model is a RUM [10]? The efficacy of the PCMC model suggests exploring other effective parameterizations for the pairwise probability P, including developing methods for inferring contractible partitions. There are also open computational questions, such as streamlining the likelihood maximization using gradients of the implicit function definitions. Very recently, learning results for nested MNL models have shown favorable query complexity under an oracle model [2], and a comparison of our PCMC model with these approaches to learning nested MNL models is important future work.

Acknowledgements. This work was supported in part by a David Morgenthaler II Faculty Fellowship and a Dantzig–Lieberman Operations Research Fellowship.

A Additional data from empirical analysis

A.1 Count data

Here we present the frequency with which pairs of alternatives appear in the same choice set. Matrices A_{work} and A_{shop} have as their (i, j) entry the number of choice sets which contained both i and j in the `SFwork` and `SFshop` datasets respectively.

$$A_{\text{work}} = \begin{bmatrix} - & 1323 & 4755 & 3729 & 1658 & 4755 \\ 1323 & - & 1479 & 1395 & 797 & 1479 \\ 4755 & 1479 & - & 4003 & 1738 & 5029 \\ 3729 & 1395 & 4003 & - & 1611 & 4003 \\ 1658 & 797 & 1738 & 1611 & - & 1738 \\ 4755 & 1479 & 5029 & 4003 & 1738 & - \end{bmatrix}$$

$$A_{\text{shop}} = \begin{bmatrix} - & 3075 & 3075 & 3075 & 1844 & 3075 & 2069 & 2916 \\ 3075 & - & 3157 & 3157 & 1908 & 3157 & 2136 & 2997 \\ 3075 & 3157 & - & 3157 & 1908 & 3157 & 2136 & 2997 \\ 3075 & 3157 & 3157 & - & 1908 & 3157 & 2136 & 2997 \\ 1844 & 1908 & 1908 & 1908 & - & 1908 & 1908 & 1876 \\ 3075 & 3157 & 3157 & 3157 & 1908 & - & 2136 & 2997 \\ 2069 & 2136 & 2136 & 2136 & 1908 & 2136 & - & 2094 \\ 2916 & 2997 & 2997 & 2997 & 1876 & 2997 & 2094 & - \end{bmatrix}$$

The relative frequencies of choice sets of different sizes in the two datasets are given in Table 1.

Table 1:

$\lvert S\rvert$ =	3	4	5	6	7	8
`SFwork`	948	1918	1461	702	-	-
`SFshop`	0	1	131	902	311	1812

A.2 Inferred pairwise probability matrices

Here we give numerical values of $\hat{P}$ inferred from the full `SFwork` and `SFshop` datasets, corresponding with the heatmaps of $\hat{P}_{\text{work}}$ and $\hat{P}_{\text{shop}}$ appearing in figure 2:

$$\hat{P}_{\text{shop}} = \begin{bmatrix} - & 0.938 & 0.875 & 0.926 & 1.000 & 0.920 \\ 0.062 & - & 0.559 & 0.650 & 0.324 & 0.998 \\ 0.125 & 0.441 & - & 0.230 & 1.000 & 0.864 \\ 0.074 & 0.350 & 0.770 & - & 0.392 & 0.899 \\ 0.000 & 0.676 & 0.000 & 0.608 & - & 0.612 \\ 0.080 & 0.002 & 0.136 & 0.101 & 0.388 & - \end{bmatrix}$$

$$\hat{P}_{\text{work}} = \begin{bmatrix} - & 0.224 & 0.740 & 1.000 & 1.000 & 1.000 & 1.000 & 1.000 \\ 0.776 & - & 0.000 & 0.312 & 0.927 & 0.532 & 1.000 & 1.000 \\ 0.260 & 1.000 & - & 0.370 & 0.617 & 0.991 & 0.565 & 1.000 \\ 0.000 & 0.688 & 0.630 & - & 0.172 & 0.593 & 1.000 & 0.656 \\ 0.000 & 0.073 & 0.383 & 0.828 & - & 1.000 & 1.000 & 0.972 \\ 0.000 & 0.468 & 0.009 & 0.407 & 0.000 & - & 1.000 & 0.617 \\ 0.000 & 0.000 & 0.435 & 0.000 & 0.000 & 0.000 & - & 0.000 \\ 0.000 & 0.000 & 0.000 & 0.344 & 0.028 & 0.383 & 1.000 & - \end{bmatrix}$$

Notice that many of the inferred pairwise parameters have values such that $p_{ij} = 1.0$ and $p_{ji} = 0.0$. These learned parameters do not indicate that i was never chosen from any sets S with i and j both in S, it merely indicates that j "passes" a maximum amount of probability mass to i in the CTMC model predicting choices on sets including both.

References

[1] E. Adams and S. Messick. An axiomatic formulation and generalization of successive intervals scaling. *Psychometrika*, 23(4):355–368, 1958.

[2] A. R. Benson, R. Kumar, and A. Tomkins. On the relevance of irrelevant alternatives. In *WWW*, 2016.

[3] J. Blanchet, G. Gallego, and V. Goyal. A markov chain approximation to choice modeling. In *EC*, 2013.

[4] H. D. Block and J. Marschak. Random orderings and stochastic theories of responses. *Contributions to Probability and Statistics*, 2:97–132, 1960.

[5] A. Börsch-Supan. On the compatibility of nested logit models with utility maximization. *Journal of Econometrics*, 43(3):373–388, 1990.

[6] J H Boyd and R E Mellman. The effect of fuel economy standards on the us automotive market: an hedonic demand analysis. *Transportation Research Part A: General*, 14(5-6):367–378, 1980.

[7] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs the method of paired comparisons. *Biometrika*, 39(3-4):324–345, 1952.

[8] S. Chen and T. Joachims. Modeling intransitivity in matchup and comparison data. In *WSDM*, 2016.

[9] G. Debreu. Review of individual choice behavior: A theoretical analysis. *American Economic Review*, 1960.

[10] J.-C. Falmagne. A representation theorem for finite random scale systems. *J. Math. Psych.*, 18(1):52–72, 1978.

[11] P. E. Green, A. M. Krieger, and Y. Wind. Thirty years of conjoint analysis: Reflections and prospects. *Interfaces*, 31(3_supplement):S56–S73, 2001.

[12] Frank Harary and Leo Moser. The theory of round robin tournaments. *The American Mathematical Monthly*, 73(3):231–246, 1966.

[13] J. Huber, J. W. Payne, and C. Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, pages 90–98, 1982.

[14] S. Ieong, N. Mishra, and O. Sheffet. Predicting preference flips in commerce search. In *ICML*, 2012.

[15] J. Kleinberg and M. Raghu. Team performance with test scores. In *EC*, pages 511–528, 2015.

[16] R. Kohli and K. Jedidi. Error theory for elimination by aspects. *Operations Research*, 63(3):512–526, 2015.

[17] F. S Koppelman and C. Bhat. A self instructing course in mode choice modeling: multinomial and nested logit models. *US Department of Transportation, Federal Transit Administration*, 31, 2006.

[18] R. D. Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, 1959.

[19] R. D. Luce. The choice axiom after twenty years. *J. Math. Psych.*, 15(3):215–233, 1977.

[20] C. F. Manski. The structure of random utility models. *Theory and Decision*, 8(3):229–254, 1977.

[21] L. Maystre and M. Grossglauser. Fast and accurate inference of plackett–luce models. In *NIPS*, 2015.

[22] D. McFadden. Econometric models for probabilistic choice among products. *Journal of Business*, pages S13–S29, 1980.

[23] D. McFadden, K. Train, et al. Mixed mnl models for discrete response. *Journal of Applied Econometrics*, 15(5):447–470, 2000.

[24] S. Negahban, S. Oh, and D. Shah. Rank centrality: Ranking from pair-wise comparisons. *arXiv preprint arXiv:1209.1688v4*, 2015.

[25] J. Nocedal and S. J. Wright. Numerical optimization. 2006.

[26] I. Simonson and A. Tversky. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29(3):281, 1992.

[27] L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273, 1927.

[28] J. S. Trueblood, S. D. Brown, A. Heathcote, and J. R. Busemeyer. Not just for consumers context effects are fundamental to decision making. *Psychological Science*, 24(6):901–908, 2013.

[29] A. Tversky. Elimination by aspects: A theory of choice. *Psychological Review*, 79(4):281, 1972.

[30] J. I. Yellott. The relationship between luce's choice axiom, thurstone's theory of comparative judgment, and the double exponential distribution. *J. Math. Psych.*, 15(2):109–144, 1977.